









Proceeding Paper

A Water Futures Approach on Water Demand Forecasting with Online Ensemble Learning [†]

Dennis Zanutto ^{1,2,*} , Christos Michalopoulos ^{1,3} , Georgios-Alexandros Chatzistefanou ^{1,4}, Lydia Vamvakeridou-Lyroudia ^{1,4}, Lydia Tsiami ^{1,3} , Konstantinos Glynis ^{1,5}, Panagiotis Samartzis ⁶, Luca Hermes ⁷, Fabian Hinder ⁷, Jonas Vaquet ⁷, Valerie Vaquet ^{7,8} , Demetrios Eliades ⁸ , Marios Polycarpou ⁸, Phoebe Koundouri ^{9,10,11,12,13} , Barbara Hammer ⁷  and Dragan Savić ^{1,4} 

- ¹ KWR Water Research Institute, 3433 PE Nieuwegein, The Netherlands; christos.michalopoulos@kwrwater.nl (C.M.); alex.chatzistefanou@kwrwater.nl (G.-A.C.); lydia.vamvakeridou-lyroudia@kwrwater.nl (L.V.-L.); lydia.tsiami@kwrwater.nl (L.T.); konstantinos.glynis@kwrwater.nl (K.G.); dragan.savic@kwrwater.nl (D.S.)
 - ² Dipartimento di Elettronica, Informazione e Bioingegneria, Politecnico di Milano, 20133 Milano, Italy
 - ³ Department of Water Resources and Environmental Engineering, National Technical University of Athens, 15773 Athens, Greece
 - ⁴ Centre for Water Systems, University of Exeter, Exeter EX4 4QD, UK
 - ⁵ Department of Water Management, TU Delft, 2628 CN Delft, The Netherlands
 - ⁶ Department of Economics, University of Macedonia, 54636 Thessaloniki, Greece; psamartzis@uom.edu.gr
 - ⁷ Machine Learning Group, CITEC, Bielefeld University, 33619 Bielefeld, Germany; lhermes@techfak.uni-bielefeld.de (L.H.); fhinder@techfak.uni-bielefeld.de (F.H.); jvaquet@techfak.uni-bielefeld.de (J.V.); vvaquet@techfak.uni-bielefeld.de (V.V.); bhammer@techfak.uni-bielefeld.de (B.H.)
 - ⁸ KIOS Research and Innovation Center of Excellence, University of Cyprus, Nicosia 2109, Cyprus; eliaades.demetrios@ucy.ac.cy (D.E.); polycarpou.marios@ucy.ac.cy (M.P.)
 - ⁹ Department of International & European Economic Studies, Athens University of Economics and Business, 10434 Athens, Greece; pkoundouri@aueb.gr
 - ¹⁰ Department of Technology, Management and Economics, Technical University of Denmark, 2800 Kongens Lyngby, Denmark
 - ¹¹ ATHENA Research Center, 26504 Patras, Greece
 - ¹² UN SDSN Global Climate Hub, 10434 Athens, Greece
 - ¹³ UN SDSN Europe, 75009 Paris, France
- * Correspondence: dennis.zanutto@polimi.it
- [†] Presented at the 3rd International Joint Conference on Water Distribution Systems Analysis & Computing and Control for the Water Industry (WDSA/CCWI 2024), Ferrara, Italy, 1–4 July 2024.



Citation: Zanutto, D.; Michalopoulos, C.; Chatzistefanou, G.-A.; Vamvakeridou-Lyroudia, L.; Tsiami, L.; Glynis, K.; Samartzis, P.; Hermes, L.; Hinder, F.; Vaquet, J.; et al. A Water Futures Approach on Water Demand Forecasting with Online Ensemble Learning. *Eng. Proc.* **2024**, *69*, 60. <https://doi.org/10.3390/engproc2024069060>

Academic Editors: Stefano Alvisi, Marco Franchini, Valentina Marsili and Filippo Mazzoni

Published: 4 September 2024



Copyright: © 2024 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

Abstract: This study presents a collaborative framework developed by the Water Futures team of researchers for the “Battle of the Water Demand Forecasting” challenge at the 3rd International WDSA-CCWI Joint Conference. The framework integrates an ensemble of machine learning forecasting models into a deterministic outcome consistent with the competition formulation. The water demand trajectory over a week exhibits complex overlapping patterns and non-linear dependencies to multiple features and time-dependent events that a single model cannot accurately predict. As such, the reconciled forecast from an ensemble of models exceeds the performance of the individual ones and exhibits higher stability across the weeks of the year and district metered areas considered.

Keywords: machine learning; ensemble forecasting; water demand forecasting; district metered areas; deep learning

1. Introduction

Accurately forecasting short-term water demand enables water utilities to make informed operational decisions and ensure an efficient and continuous water supply. A wide array of methods and techniques exist in the literature to model urban water demand [1].

The water consumption patterns are influenced by multiple factors, such as the area characteristics (e.g., city center vs. rural), seasonality, weather and time-dependent events (e.g., water leaks), and societal habits (e.g., holidays), and no single framework has yet emerged as universally superior. Moreover, water utilities' specific needs (e.g., resolution, lead time) can only be satisfactorily addressed by tailored models.

Aware of the limitations of any single forecasting methodology, the Water Futures (<https://waterfutures.eu/> (accessed on 30 March 2024)) approach is distinguished by its integration of multiple models developed by experts from diverse research fields (civil engineering, economics, machine learning, and control). This collaborative effort is tailored to address the practical challenges of the Battle of the Water Demand Forecasting (BWDF), hosted during the 3rd International WDSA-CCWI Joint Conference.

2. Context and Data

The goal of the Battle of the Water Demand Forecasting (BWDF) is to produce the best water demand forecast for 10 district metered areas (DMAs) of an undisclosed city in Northern Italy, with an hourly resolution over a one-week horizon. The forecasts' quality is evaluated independently for each DMA on three Performance Indicators (PIs): (i) Mean Absolute Error in the first 24 h, (ii) Maximum Absolute Error in the first 24 h, (iii) and MAE for the remainder of the week. The evaluation is conducted across four distinct weeks; at every assessment, T weeks of raw measurements on the net DMA water consumption and $T + 1$ of weather observations (rain, temperature, humidity, and windspeed) with an hourly resolution are disclosed, and the water demand of week $T + 1$ has to be forecasted (the additional weather week can be trusted as perfect information).

3. Methods

The Water Futures framework produces a deterministic water demand forecast combining an ensemble of models of different natures. The core concept revolves around the idea that for each DMA and part of the week (first day vs. rest of the week), there is a model, or a combination of, that best performs on forecasting that signal. The strategy to select and combine the models can be static or dynamic (online). To ensure impartiality, a batch of testing data is extracted upfront. The models' performance on it guides the final model selection for the upcoming week's forecast. To accommodate the substantial differences between the models, the input pre-processing is integrated as a component of the model itself. This section highlights the key attributes of the more sophisticated models; for the full list, see [2].

3.1. Gradient Boosting Models

LightGBM [3] and XGBoost [4] are two algorithms that implement gradient boosting, usually employing regression trees as weak learners. They sequentially introduce new learners to correct prior errors, systematically minimizing a loss function to enhance prediction accuracy over iterations.

LightGBM excels in processing large datasets and grows the decision trees leaf wise, i.e., it focuses on the nodes that lead to the most significant reduction in loss, enhancing both computational efficiency and model accuracy. In contrast, XGBoost employs a level-wise growth strategy, exploring all possible splits at each level, which might reveal different data patterns. Both models share similar input features for training, including rolling statistics, lagged values, seasonality indicators, weather data, and outputs from simpler forecasting models.

These two algorithms are particularly effective for complex time series forecasting, like water demand, offering advantages such as interpretability, inherent handling of missing values and categorical variables, and built-in overfitting prevention.

3.2. Wavenet

WaveNet [5] is a deep convolutional neural network consisting of two main components: (i) a stack of one-dimensional convolutions and (ii) a multilayer perceptron (MLP). The convolutional part employs causal convolutions with a receptive field of 32 h. This single-step prediction is extended to a multi-step forecast of one week in an autoregressive fashion.

To counteract overfitting and enhance the model's robustness towards missing data, we employ dropout regularization [6]. This regularization technique is applied uniformly throughout the model; instead of masking individual input features, as commonly performed, it masks random temporal segments of different lengths of the input time series.

3.3. Ensemble Reconciliation Strategies

Ensemble reconciliation strategies (ERSs) aim to refine multiple model forecasts into an optimal deterministic forecast by evaluating models based on past performance and ground truth data. These strategies select a subset of top-performing models, re-train them on the full dataset, and combine their forecasts for the upcoming week. ERSs can be categorized into two types: selecting the best model from the past w weeks or combining forecasts from the top models. While more dynamic strategies could be considered, such as choosing models based on similarity to the upcoming week, they require more comprehensive datasets.

4. Experimental Settings

Most of the decisions and settings have been tailored to maximize the performance at the BWDF. However, the Water Futures framework could be adapted to other contexts and problem formulations. The testing dataset comprises the four weeks before the one to be forecasted, assuming that the preceding month is a reliable indicator, *on average*, of a model's prospective performance in the coming week. A suitable evaluation scheme must (i) be suited to the time series properties of the data, (ii) avoid data leakage for an appropriate comparison between models, and (iii) be feasible with the limited dataset available.

Thus, we leveraged the *test-then-train* scheme, which is frequently used in online or stream learning by the machine learning community [7,8]. Splitting data into independent train and test sets is not feasible when the data distribution changes over time. Thus, we employ data up to time t for training and testing on the next observed sample ($t + 1$). In the following step, the model is updated with the newly acquired sample ($t + 1$) and tested on $t + 2$, and this cycle continues. This yields an unbiased estimate of the model's performance throughout the considered stream. We consider a test batch one week at a time, according to the formulation of the challenge. During model development, the process is carried out for 52 weeks of the training dataset to avoid bias towards any season.

Multiple trainings with different initialization seeds are run for the non-deterministic models. Moreover, different versions of LGBM varying in the input dataset are trained.

Several strategies have been tested, but the chosen ERS averages the forecasts of the five best models.

5. Results

The gradient boosting family of models is the best performing, with WaveNet following, but it still shows more skills than the benchmark (Table 1). The ERS combines these five models in a single deterministic forecast that outperforms the individuals, proving that the ensemble methodology adopted here is beneficial.

Table 1. Average performances (L/s) of the five best models, the benchmark (rolling average), and the chosen ensemble reconciliation strategy (average of the five best models) during training.

PI	Rolling Average	LGBM Robust	LGBM Simple	LGBM SimpleV2	XGBOOST	WAVENET	ERS
Mean AE-24h	1.286	1.080	1.078	1.099	1.102	1.226	1.044
Max AE-24h2	4.040	3.550	3.530	3.596	3.639	3.912	3.460
Mean AE	1.163	1.069	1.067	1.073	1.083	1.197	1.016

6. Conclusions

In this work, we introduced a collaborative framework that combines multiple machine learning forecasting models and applied it to the “Battle of the Water Demand Forecasting (BWDF)”, an urban water demand forecasting challenge hosted during the 3rd International WDSA-CCWI Joint Conference. Our findings indicate that the gradient boosting family (e.g., LGBM) is particularly effective for predicting water demand across most DMAs and throughout the week. Despite their effectiveness, no single model can fully capture these complex patterns. The ensemble reconciliation strategy (ERS) enhanced performances (the ensemble consistently beats the individual models) and stability across the evaluated weeks and the DMAs.

Author Contributions: Formal analysis, software, and writing—original draft preparation, D.Z., C.M., P.S., L.H., J.V., V.V. and D.E.; methodology and writing—review and editing D.Z., C.M., L.T., K.G., G.-A.C., L.V.-L., P.S., L.H., F.H., J.V., V.V. and D.E.; supervision and funding acquisition, M.P., P.K., B.H. and D.S. All authors have read and agreed to the published version of the manuscript.

Funding: This project has received funding from the European Research Council (ERC) under the ERC Synergy Grant Water-Futures (grant agreement No. 951424).

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The code is available on Github (<https://github.com/WaterFutures/BoN2024/tree/v1.0.0> accessed on 30 March 2024) [2]. The results are open access on Zenodo [9].

Acknowledgments: We extend our heartfelt gratitude to all KWR colleagues who engaged in many insightful discussions and to the academic supervisor of the KWR Ph.D. students.

Conflicts of Interest: The authors declare no conflicts of interest.

References

- House-Peters, L.A.; Chang, H. Urban Water Demand Modeling: Review of Concepts, Methods, and Organizing Principles. *Water Resour. Res.* **2011**, *47*. [[CrossRef](#)]
- Zanutto, D.; Michalopoulos, C.; Samartzis, P.; Hermes, L.; Vaquet, J.; Vaquet, V.; Eliades, D. *WaterFutures/BoN2024: BWDF Code, Models and Dashboard, Version v5.0.0*; Zenodo: Geneva, Switzerland, 2024. [[CrossRef](#)]
- Ke, G.; Meng, Q.; Finley, T.; Wang, T.; Chen, W.; Ma, W.; Ye, Q.; Liu, T.-Y. LightGBM: A Highly Efficient Gradient Boosting Decision Tree. In Proceedings of the Advances in Neural Information Processing Systems, Long Beach, CA, USA, 4–9 December 2017; Curran Associates, Inc.: New York, NY, USA, 2017; Volume 30.
- Chen, T.; Guestrin, C. XGBoost: A Scalable Tree Boosting System. In Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, San Francisco, CA, USA, 13–17 August 2016; Association for Computing Machinery: New York, NY, USA, 2016; pp. 785–794.
- van den Oord, A.; Dieleman, S.; Zen, H.; Simonyan, K.; Vinyals, O.; Graves, A.; Kalchbrenner, N.; Senior, A.; Kavukcuoglu, K. WaveNet: A Generative Model for Raw Audio 2016. *arXiv* **2016**, arXiv:1609.03499.
- Srivastava, N.; Hinton, G.; Krizhevsky, A.; Sutskever, I.; Salakhutdinov, R. Dropout: A Simple Way to Prevent Neural Networks from Overfitting. *J. Mach. Learn. Res.* **2014**, *15*, 1929–1958.
- Ditzler, G.; Roveri, M.; Alippi, C.; Polikar, R. Learning in Nonstationary Environments: A Survey. *IEEE Comput. Intell. Mag.* **2015**, *10*, 12–25. [[CrossRef](#)]

8. Bifet, A.; Holmes, G.; Pfahringer, B.; Gavalda, R. Improving Adaptive Bagging Methods for Evolving Data Streams. In *Advances in Machine Learning*; Zhou, Z.-H., Washio, T., Eds.; Springer: Berlin/Heidelberg, Germany, 2009; pp. 23–37.
9. Zanutto, D.; Michalopoulos, C.; Samartzis, P.; Hermes, L.; Vaquet, J.; Vaquet, V.; Eliades, D. *WaterFutures/BoN2024: BWDF Input Data and Results, Version v4.0.0*; Zenodo: Geneva, Switzerland, 2024. [[CrossRef](#)]

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.